

Wiley, Rein & Fielding

1776 K Street, N.W.
Washington, D.C. 20006
(202) 719-7000

Peter D. Ross
(202) 719-4232
pross@wrf.com

Fax: (202) 719-7049
www.wrf.com

June 29, 2000

RECEIVED
JUN 29 2000
FEDERAL COMMUNICATIONS COMMISSION
OFFICE OF THE SECRETARY

By Hand

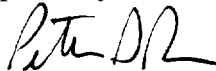
Ms. Magalie R. Salas
Secretary
Federal Communication Commission
445 12th Street, S.W. - The Portals
Washington, D.C. 20554

Re: Response to June 9, 2000 Request for Further Information
Applications of America Online, Inc. and Time Warner Inc.
for Transfers of Control
CS Docket No. 00-30

Dear Ms. Salas:

Transmitted herewith on behalf of America Online, Inc. ("AOL") are additional materials, Bates Nos. AOL 000345 through AOL 000365, referred to in our June 26, 2000 narrative response to Question 1.13 of the June 9, 2000 letter from Ms. To-Quyen Troung, Associate Chief of the Cable Services Bureau, requesting certain documents and information. Please do not hesitate to contact me should you have any questions regarding this letter or the documents produced herewith.

Respectfully submitted,



Peter D. Ross
Counsel for America Online, Inc.

cc w/ Attachment: Royce Dickens, Cable Services Bureau
Linda Senecal, Cable Services Bureau
Darryl Cooper, Cable Services Bureau
James Bird, Office of General Counsel

No. of Copies rec'd 0
List ABCDE



▶ AOL Mail ▶ My AOL.COM ▶ People/Chat ▶ Search ▶ Shop ▶ Web Centers ▶ Try AOL FREE!

Open IM Architecture Design

America Online is committed to extending the benefits of instant messaging technology to as many consumers as possible. We give our AOL Instant Messenger client software away for free to anyone on the Internet, and have entered into more than a dozen royalty-free license agreements to allow industry leaders, including Lotus, Lycos, Earthlink and other ISPs, to distribute the software to their customers.

Over the past two years, we have participated in industry discussions through the IETF about how to achieve the goal of interoperability among instant messaging networks. At the same time, we have resisted efforts by our competitors to impose a "quick fix" system that would jeopardize our members' privacy and security. Both of these decisions have been guided by two bedrock commitments: To provide consumers with a secure, private, and convenient online experience; and to help build a medium we can all be proud of.

In response to your call for instant messaging ideas to be submitted by today, we are reaffirming our consistent commitment to interoperability with the release of our proposed architectural design for a worldwide instant messaging system. We think that this submission represents a significant first step toward developing more detailed protocols for implementing the kind of full interoperability that we all would like to see.

The IETF knows better than anyone how complex and difficult an engineering challenge is posed by developing standards that protect the consumer. Ultimately, our first commitment and biggest concern in this process continues to be protecting the privacy and security of consumers. Once protocols are published, they will be used by hackers and spammers as a roadmap to plan their attacks. We believe that it is critically important not to release such proposals until we are certain that the security precautions in them are sufficient to protect consumers. If we do not move deliberately through this process, we risk undermining our efforts to serve consumers well. Nobody wants to see instant messaging interoperability made useless by a barrage of offensive spam, attempts to defraud, and virus proliferation.

The design we have submitted is a major advance toward developing a system that forestalls these threats with a server-to-server approach to interoperability that would work in a manner similar to Internet e-mail. We believe that this approach protects the user's privacy, security and ease-of-use as well as promoting continued long-term competition and innovation in the industry and providing the greatest degree of scalability. Among the advantages of this system are:

- **Full Interoperability:** The server-to-server architecture of the AOL design allows users of any two IM networks that use the same protocols to communicate with one another at any time;
- **Privacy and Security:** Under this architecture, users would need to be registered with only one instant messaging system, and would not be required to share passwords, log-in IDs, or other confidential information with anyone outside the network they choose. In addition, the design includes requirements that IM data can't be easily spoofed or replayed by a third party; that messages can't easily be intercepted or hijacked; and that more advanced security measures such as end-to-end encryption or signing can be layered on top of initial implementations. Finally, it requires that individual networks be allowed to use firewalls or other precautions to ensure the highest possible degree of security;
- **Scalability:** This design would enable the development of any number of instant messaging systems, from large networks like AOL to individual families with their own servers;
- **Independence:** No government or other central authority would be required to administer the system; and
- **User Name Consistency:** Users of instant messaging networks would be able to keep their existing screen names or addresses, even if different users had identical names on different networks.

In addition to openly publishing our design, we are publishing this document as an Internet Draft and submitting it to the IETF for consideration as an informational RFC, in the hope that it will serve as a helpful guide to the standards-development process, and that other participants in the instant

as a helpful guide to the standards-development process, and that other participants in the instant messaging working group will offer useful comments on this design as we work to develop protocols for its implementation.

[Text of Internet Draft - The Open IM Architecture](#)
[IMX Architecture Diagram](#)

Other AOL Sites

[AOL Hometown](#)
[AOL Instant Messenger](#)
[AOL Affiliate Network](#)

[AOL Mail](#)
[New to Chat?](#)

[Find a Chat](#)
[People & Chat Directory](#)

[Download AOL 5.0](#)
[Love@AOL](#)

Copyright © 2000 America Online, Inc.
All rights reserved. [Legal Notices](#)
[Privacy Policy](#)
[Try AOL 5.0](#)

INTERNET-DRAFT
<draft-aol-umx-00.txt>

E. Aoki
A. Wick
AOL

Expires: December 15, 2000

June 15, 2000

The IMX Architecture
Interoperability with America Online's Instant Messaging Services

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This memo provides information for the Internet community.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

The ability to exchange instant messages and presence information provides users with a powerful mechanism for communicating in real-time.

This document outlines an architecture for interoperability among Instant Messaging Systems which allows disparate systems to exchange messages and presence information while being relatively easy to implement and maintaining a high standard of security and scalability.

Table of Contents

1. Introduction	2
2. Requirements	3
3. Terminology	4
4. Architecture	6
4.1 Servers and Gateways	6

[Page 1]

Aoki & Wick

INTERNET-DRAFT

IMX Architecture

June 15, 2000

4.2 Namespace and Addressing	7
4.2.1 Address identifiers	7
4.2.2 Domains	7
4.2.3 Server Discovery	7

- 4.3 Connection Management 8
- 4.3.1 Originating Connections 8
- 4.3.2 Accepting Connections 8
- 4.3.3 Relays 9
- 4.4 Protocol Considerations 9
- 4.5 Additional Protocol Considerations for Presence Information 10
- 4.6 Additional Protocol Considerations for Attributes ... 10
- 4.7 Additional Protocol Considerations for Instant Message Information 10
- 5. Instant Message Format Considerations 11
- 6. Security Considerations 11
- 6.1 Objectives 11
- 6.2 Assumptions 12
- 6.2.1 Client Authentication 12
- 6.2.2 Scope 12
- 6.2.2.1 Types of Attacks Within the Scope 13
- 6.2.2.2 Types of Attacks Outside of Scope 13
- 6.3 A Trust Model for Server to Server communications.... 14
- 6.3.1 The Dial-Back Mechanism 14
- 6.3.2 Enhanced Security 15
- 6.4 SPAM 16
- 7. Authors' Addresses 16
- 8. Additional Documents 17
- 9. Acknowledgements 17
- 10. References 17
- A. Appendix - Performance Against Objectives 17
- A.1 Scalability and Efficiency 18
- A.2 Ease of Implementation 19
- A.2.3 Reliability 19
- A.2.4 Security 19

1. Introduction

Today's instant messaging systems are typically comprised of a client, through which the end-user interacts, and servers which relay information between compatible clients. Tight integration between clients and servers allows instant messaging services to provide a secure, reliable channel through which authentication, presence, and messaging information is passed between users and the service.

As the number of instant messaging providers has grown, there is increased interest in enabling IM users to exchange presence and messaging information not only with users on their system, but with those on other systems as well. Some vendors have responded by creating "multi-headed clients," clients which can simultaneously communicate with servers on disparate instant messaging systems.

[Page 2]

Aoki & Wick

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Such clients achieve interoperability at a high price, however. Since each service has its own feature set, clients may advertise features that do not work across systems. Since each service implements its own security model, multi-headed clients must often resort to mechanisms that circumvent security or require the user to provide passwords to third parties. Inconsistent terms of service also make it difficult to enforce anti-spam measures or encourage equitable resource sharing. And vendors are forced to constantly upgrade clients to keep up with changes in features and services across the instant messaging universe.

An alternative approach is to provide a mechanism for the services

themselves to interoperate in a peering arrangement, much like the Internet mail system works today. In such a system, the interaction between instant messaging clients and their associated servers would remain much as it is today, but servers could communicate with other servers to exchange presence information, messages, or other data. This approach helps to preserve existing models for security and allows Instant Messaging Service providers to manage client authentication, service policy, and privacy.

This document describes how America Online intends to develop such an architecture to allow its services to discover other servers (and be discovered by other servers), exchange data, and ensure security. It also describes implications that any architecture for interoperability may have for the spread of unsolicited instant messaging (spam).

2. Requirements

The authors set out to design a system that would be flexible, yet practical to implement. In that vein, many of our design goals, listed below, are the same as or similar to those specified in [RFC 2779], but with additional consideration paid to implementation issues.

The primary requirements were to design a system which:

- 1) is scalable to hundreds of millions of instant messaging users.
- 2) is scalable to hundreds of thousands (or perhaps millions) of individual instant messaging systems and domains, so every company or ISP could have its own instant messaging system.
- 3) allows existing instant messaging implementations to manage their own user namespace.
- 4) is self managed (like the Internet mail system) such that new servers and new systems could be added without administration

Aoki & Wick

[Page 3]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

by some central authority and without manual administration by other service providers.

- 5) initially supports at least five main end-user features:

- Requesting/Renewing/Canceling presence subscriptions
- Sending/Receiving presence notifications (in response to a subscription)
- Routing and delivery of instant messages
- Retrieval of named user attributes (at minimum an "alias" and current presence state)
- Retrieval of named domain attributes

- 6) allows traffic between users in a single instant messaging system to stay within that system.

- 7) makes it possible to implement interoperability between instant messaging systems by adding gateways to existing systems without rearchitecting existing core systems.

- 8) is extensible, so new features can be added incrementally without requiring redesign and while allowing for backwards

AOL 000349

compatibility.

9) has a well thought-through security strategy such that:

- Messaging data or state can't be easily spoofed or replayed by a third party
- Messaging data or state can't be easily intercepted, hijacked, or stolen by a third party
- More advanced security measures such as end-to-end encryption or signing can be layered on top of the initial implementation, but are not required in the initial implementation.

10) easily supports clients that are inside of a common company firewall (e.g. incoming connections are often refused).

11) easily supports international usage.

12) leverages existing standards in as many places as practical.

3. Terminology

[RFC 2778] specifies a common terminology for the discussion of Internet Messaging and Presence Protocol information; however, that document defines terms which are considerably more granular than are required by this document. Accordingly, this document uses the following terms:

Aoki & Wick

[Page 4]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Aggregator - A special kind of Gateway, which services multiple domains and routes messages between other IMX Servers and Servers which service a particular domain. The protocol between the Aggregator and each Server is arbitrary. An example of an Aggregator might be a Gateway which acts as a front-end to multiple, privately-labeled Instant Messaging Services.

Attributes - metadata about an End User, such as a nickname or alias; or about a domain, such as a timeout value. Attributes consist of a key, which is scoped at a domain, and a value. It may be desirable to have certain attribute keys which are global to the IMX architecture and interpreted identically by all participating services.

Data - any of Instant Messages, Attributes, Notifications, Subscriptions, or requests for these that are exchanged between Servers.

End User - a human or other entity whose presence information is reflected by the service and who can send and receive Instant Messages through an Instant Messaging Service.

Instant Message - a short, real-time or near-real-time message which is sent between Instant Messaging clients. While this document does not prescribe a definition for "short," the intent is to prevent streams of arbitrary length from being sent as Instant Messages. This is consistent with the definition of an INSTANT MESSAGE in [RFC 2778].

Instant Messaging Client (or Client) - a User Agent which provides

an End User with the ability to initiate and receive Instant Messages, and request Subscriptions and receive Notifications for Presence Information. This term is used in this document to encompass a SENDER, INSTANT INBOX, PRESENTITY, and WATCHER in [RFC 2778], and is roughly consistent with the generic description of an "instant messenger" in section 2.7 of that document.

Instant Messaging Gateway (or Gateway) - a special type of Instant Messaging Server which does not communicate directly with Clients but sits between Servers which service a particular domain and the world of other IMX servers. Gateways act essentially as routers, potentially performing protocol translation between an Instant Messaging Service's local protocols and the IMX protocol. An example of a Gateway would be a Server which routes messages to an already existing, private Instant Messaging Service.

Instant Messaging Server (or Server) - an entity which maintains Presence Subscriptions for and delivers Instant

Aoki & Wick

[Page 5]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Messages on behalf of End Users of a given Instant Messaging Service. Typically, Instant Messaging Clients connect to a Server in order to send and receive Instant Messages, Attributes, and Presence Information.

Instant Messaging Service - a collection of Instant Messaging Servers and their associated clients which together make up an integrated service offering. Instant Messaging Services administer their own namespace of End User names and are generally associated with one or more domains they serve.

Notification - a message sent from a Server to a Client to indicate a change in the Presence Information for a given End User. Notifications are sent only to Clients who have previously created a Subscription to receive such information for a particular End User. (Presence Information may be retrieved in real-time without a Subscription by making a request for a presence Attribute). This term is essentially identical to the corresponding definition in [RFC 2778].

Presence Information - information regarding the state of a particular user. Examples might be online, offline, away, busy, etc. Specific formats for the conveyance of Presence Information are not specified here.

SPAM - unwanted (and typically unsolicited) Instant Messages.

Subscription - Information kept by an Instant Messaging Server which maintains a Client's desire to be Notified when an End User's Presence Information changes. Subscriptions are entered by Instant Messaging Clients on behalf of End Users, and time out if not renewed. This term is essentially identical to the corresponding definition in [RFC 2778].

4. Architecture

The basic architecture behind IMX (Instant Messaging eXchange) is one of multiple independent Instant Messaging Services (as exist today), which can interoperate at the Server-to-Server, Server-to-Gateway, or

Gateway-to-Gateway level.

4.1 Servers and Gateways

Each Instant Messaging Service exchanges Data with other Services through well known Servers or Gateways. For the purpose of this document, Servers and Gateways are distinguished in that Servers can be the same physical or logical entities that service Clients, whereas Gateways essentially relay information between Servers within a Service and those outside of it. Absent this distinction, the terms Server and Gateway can be used interchangeably.

[Page 6]

Aoki & Wick

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Each Server is said to "service", "be responsible for", "act on behalf of," or be "in" a domain, which is a shortcut to saying that a Server listens for requests intended for a given Instant Messaging Service (identified by a domain), and provides answers on behalf of the End Users who belong to that Service. In this way, the model is very similar to the way that SMTP servers exchange mail with other SMTP servers.

4.2 Namespace and Addressing

4.2.1 Address identifiers

Each End User has an identifier which can be used to uniquely address that user across various Instant Messaging Systems. Like internet mail addresses, the identifier used to identify a given End User consists of a local part and a domain, separated by at-sign (@). (The at-sign is therefore a reserved character in an Instant Messaging address).

The address used for Instant Messages may or may not be the same as an End User's email address.

4.2.2 Domains

The domain portion of the address is assumed to be an internet domain, compliant with [RFC 1034] and is used to identify the Servers responsible for that domain by the process described in Section 4.2.3.

4.2.3 Server Discovery

Servers which service a particular domain are published in DNS using an IMX record. An IMX record is a new type of DNS record that works similar to the MX record used by mail. It lists one or more Servers that accept Instant Messaging connections for that particular domain. If more than one server is listed, the originating Server picks one of the listed servers randomly (just like with MX records) and attempts to establish a connection. If the chosen server doesn't respond, the originating server may attempt to connect on one of the other listed servers.

The IMX record may also list other attributes of the connection such as a port number.

A full specification of IMX records will follow in a separate document.

4.3 Connection Management

Connections between Servers use TCP for transport. They are established on-demand by the originating Server as needed and are semi-persistent. Connections are kept open for as long as the originating host desires, as long as there is activity on the connection. Connections may be closed by the receiving Server after some period of inactivity; they would then be re-established by the originator on-demand when next needed.

Data for multiple End Users can and likely will flow across a given connection. Conversely, Data for a given user may flow over any active connection that exists between the two domains and is not guaranteed to flow over the same connection as previous requests. If a Server handles requests for more than one domain, Data for multiple domains may flow over a given connection. These connections must be authenticated multiple times, once for each domain they serve.

4.3.1 Originating Connections

A Server would establish a connection to another Server in order to request Data on behalf of an End User within the originating server's domain. This would include connections to send Instant Messages, initiate or refresh Subscriptions, send Notifications, and to request Attributes.

A Server may establish more than one connection to communicate with another Server (i.e. for load balancing or to handle increased traffic); however, in consideration of scalability, implementers should take care to limit the number of such connections to a reasonably small number. For example, it would not be appropriate for a Server to establish one connection per End User; it would be appropriate to establish one connection per Server process or Server thread. In any event, a receiving server may refuse to accept more than a given number of connections from Servers at a single domain.

Servers will also establish connections in order to perform authentication handshaking as specified in Section 6.3.1.

4.3.2 Accepting Connections

A Server would accept a connection from another Server in order to receive Data for an End User within the receiving server's domain. This would include connections to receive Instant Messages, update or establish Subscriptions, receive Notifications, and receive requests for Attributes.

Servers will also accept connections in order to perform authentication handshaking as specified in Section 6.3.1.

4.3.3 Relays

Server should only accept traffic for the domains they service. That is, this architecture explicitly disallows relaying services because of their possible unauthorized and unintended use (see SPAM, section 6.4).

Instant Messaging Services will be held accountable for traffic originating from their Servers.

4.4 Protocol Considerations

This document does not specify the details of the protocol between Servers (the IMX protocol). However, we believe that the following requirements should be considered in the implementation of the protocol:

- 1) The protocol must have the ability to exchange protocol version information and some protocol capability information. (4.4.1)
- 2) The protocol must be completely extensible without ever having to force all other servers to upgrade. The intention is that, in most cases, new types of data packets or request packets can be added, and older servers can simply ignore the data or request types that they don't understand. (4.4.2)
- 3) The protocol should include a means for either side to close the connection in an orderly fashion and have the other side know that the connection has been closed. (4.4.3)
- 4) The protocol should include some sort of redirection mechanism so that the other end of an existing connection can be told to reconnect on another IP address. (4.4.4)

This would be useful for:

- Taking an existing, active server out of service in an orderly fashion
 - Smarter load balancing beyond the capabilities of a DNS rotor.
- 5) The protocol needs the ability to retrieve named user Attributes (e.g. the user's alias, UA capabilities, maximum message length, etc...). (4.4.5)
 - 6) The protocol also needs the ability to retrieve named domain-specific Attributes. For example, the timeout for a connection or a presence subscription to expire may be a domain-specific configuration. (4.4.6)

Aoki & Wick

[Page 9]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

- 7) The protocol is completely asynchronous. The connection is simply a pipe through which requests and tagged responses flow. More than one connection may be open between a set of servers, and more than one request may be in process on a given connection. In some cases, a response may not flow back on the same connection on which it was initiated. (4.4.7)
- 8) The protocol should allow for responses to at least some requests (particularly those requesting Attributes), to return on the same connection as the request. The protocol should additionally provide enough information in the request so that a reply can be

routed to the same Server that made the request. (4.4.8)

- 9) The protocol should be binary, not text. The binary format will be records with length and type. (4.4.9)
- 10) Where not already specified by a chosen, existing standard, data will be represented in UTF-8. (4.4.10)

4.5 Additional Protocol Considerations for Presence Information

When used for Subscription or Presence Information, the following additional requirements should be considered:

- 1) Presence Subscription should expire in some pre-determined amount of time (e.g. 30 minutes) if they aren't renewed. This timeout may be different per vendor and should be published as a domain attribute. (4.5.1)
- 2) Presence Notifications may be delivered over any active connection that exists between the two domains. They are not necessarily delivered over the same connection on which the subscription was requested (that connection may not even exist anymore). (4.5.2)

4.6 Additional Protocol Considerations for Attributes

When used to retrieve Attributes for a user or domain, the following additional requirements should be considered:

- 1) Attribute information may be sent unsolicited with Instant Message or Presence Information so that it need not be explicitly requested. (4.6.1)

4.7 Additional Protocol Considerations for Instant Message Information

When used for Instant Messages, the following additional requirements should be considered:

- 1) As with the SMTP envelope, the To and From address flows with the message as part of the protocol and does not need to be parsed by Servers, Gateways, or Aggregators. (4.7.1)

Aoki & Wick

[Page 10]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

5. Instant Message Format Considerations

In order to achieve maximum flexibility, a subset of MIME [RFC 2045] should be used for the message transfer format. Messages will consist of headers and a body.

- 1) The message From and To headers are also included in the MIME message and are required. All other headers are optional. Any number of headers may exist. Headers such as "Content-type" and "Content-transfer-encoding" are used if present. "Subject" is not used. Vendors that specify additional vendor-specific headers are strongly encouraged to use the "X-" header format.
- 2) Initial implementations of the architecture should support at least two MIME types: text/plain and text/html-lite.

Text/html-lite is a strict subset of HTML with support for only a few basic formatting tags supported (such as ..., <I>...</I>, etc...).

No additional attempt is made to describe text/html-lite in this document.

- 3) Very long lines are permitted. Text/plain messages are not prewrapped to any particular column width. A paragraph is sent as one long line. It is the receiving Client's responsibility to wrap the message for proper display. It is important that all servers support this correctly because Instant Messaging user interfaces are often used in a small screen form factor where appropriate wrapping substantially increases the usability. The small screen requirement comes about either because the Client runs on a small screen device or because the user wishes to devote only a small piece of screen real estate to the Client user interface.

6. Security Considerations

6.1 Objectives

[RFC 2779] specifies in detail a series of requirements for security in an Instant Messaging protocol. In developing this architecture, the authors also considered the following objectives in order to ensure practicality of implementation:

- 1) The architecture should be straightforward to implement and administer and be freely exportable.
- 2) The system, including the relationship between arbitrary Instant Messaging Services, should be self-managed, with no centralized authority required to administer the security scheme.

[Page 11]

Aoki & Wick

INTERNET-DRAFT

IMX Architecture

June 15, 2000

- 3) The security scheme should account for a basic level of security but be easily extensible so that some Services may implement a higher degree of security.
- 4) In order to address these objectives (1-3), the intent is to authenticate connections between Servers, not individual exchanges of Data.
- 5) In keeping with the objectives of (3), it should be possible, but not required, to support end-to-end encryption and/or signing. (End-to-end encryption refers to data encrypted by the sending Client and decrypted at the recipient's Client).
- 6) In keeping with the objectives of (3), it should be possible, but not required, to support the use of advanced security measures such as SSL or certificate exchange between Servers.
- 7) User passwords should never be exchanged in a Server to Server communication.

The overarching intent behind these objectives is to provide a system which would offer substantially higher security than Internet email, yet offer a basic level of security that could be easily and rapidly implemented. At the same time, the model would be extensible enough to allow for increased security should users, vendors, or other entities demand it.

6.2 Assumptions

6.2.1 Client Authentication

Because this architecture specifies a Server to Server protocol, it does not specify or recommend any mechanism for authenticating Clients to a Server. Each Instant Messaging Service is held responsible for security between Clients and Servers, including preventing one End User from impersonating another.

In all discussions in this section, this document assumes that a given Server can be trusted when it represents that Data originates from a given End User.

6.2.2 Scope

Each individual Instant Messaging Service is responsible for making sure that a given End User is not an impersonator. Therefore, the principal problem this document tries to solve is to ensure that Data sent and received on behalf of an End User in a given domain is actually originating from and/or being sent to Servers in that domain.

Aoki & Wick

[Page 12]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Additionally, the architecture seeks to protect against a spoofing attack where one Server pretends to issue Data on behalf of users in a domain that it does not represent.

6.2.2.1 Types of Attacks Within the Scope

The following types of attacks are explicitly covered by the proposed architecture.

- 1) Breach of Presence Privacy - Spoofing an End User who is requesting Presence Information (or Attributes). This form of security becomes especially important if a service supports restrictions on who can and can not retrieve Presence Information (or Attributes) for a given End User. This is essentially the requirement 5.1.1 in [RFC 2779], where A and B are SERVERS instead of non-ADMINISTRATOR PRINCIPALS.
- 2) False Presence - Spoofing state about a given End User (e.g. if User A in a given domain is subscribed to receive Presence Information for User B in another domain, it is not possible for a User C to send Notifications with User B's Presence Information to User A). This is essentially the requirement 5.2.4 in [RFC 2779], where users A and B are authenticated ENTITIES in separate DOMAINS.
- 3) Spoofing Sender - Spoofing the From address for a given Instant Message (e.g. User C in a given domain sends an Instant Message to User B purporting to be User A). This is essentially the requirement 5.4.3 in [RFC 2779], where A and B are authenticated ENTITIES in separate DOMAINS.
- 4) Stealing Data - Hijacking Data intended for a given End User and redirecting it somewhere else. (e.g. User A sends an Instant Message to User B, but it instead is directed to User C). This is essentially requirements 5.3.3 and 5.4.6 in [RFC 2779], where A and B are authenticated ENTITIES in separate DOMAINS.

6.2.2.2 Types of Attacks Outside of Scope

For practical reasons (see Objectives, 6.1), a "standard" implementation of security as defined in this document does not protect against the following types of attacks:

- DNS Hijacking
- Man-in-the-middle attacks
- Eavesdropping or packet snooping

However, an advanced implementation of security may protect against these within the specified architecture (see section 6.3.2).

Aoki & Wick

[Page 13]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

6.3 A Trust Model for Server to Server communications

This document therefore proposes the following architecture for establishing trust between Servers:

- 1) Servers will trust DNS for all outgoing connections. If a Server establishes a connection from Domain A to Domain B, it should be relatively certain that data sent over that connection is actually going to Domain B. The only situations where that assumption would not be true would be if the entirety of Domain B were hijacked, and that is explicitly not covered by this model (see 6.2.2.2).
- 2) Incoming connections to a Server are not trusted (since they could originate anywhere). Incoming connections are authenticated by using a mechanism referred to as "dial-back" (see 6.3.1), which establishes an outbound connection to the originating Server to authenticate the Server which makes the inbound connection.
- 3) Once a connection is established and trusted, it is trusted for the life of the connection.

6.3.1 The Dial-Back Mechanism

The Dial-Back mechanism is so-named after the (now somewhat old) concept which goes something like this:

- If User A calls User B, then User A trusts that she is talking to User B because she made the call.
- If User A calls User B, User B does not trust that he is talking to User A, since anyone could call User B pretending to be User A.
- If User A calls User B and provides some secret, User B can call User A on another line and provide that same secret. This ensures that User A and User B are the same on both calls (since they have exchanged the secret), and that Users A and B trust each other (since each user originated the call on which they are providing the information). Therefore, both lines can be trusted equally.

It is fairly straightforward to extend this concept to the exchange of Data between Servers. When a Server at Domain A accepts a connection which supposedly originates from Domain B, it must first authenticate that connection as follows:

- 1) The incoming connection from Domain B (call this Connection 1) contains an initial token from Domain B. This token is purely for

Domain B's implementation convenience and can be anything Domain B wants it to be (see Implementation Note, below).

Aoki & Wick

[Page 14]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

- 2) Domain A then creates an outgoing connection to Domain B (call this Connection 2). In establishing this connection, Domain A uses no information passed in as part of the original request on Connection 1 (except, of course, the name of Domain B). That is, Domain A will use its own DNS resolution to establish its connection to Domain B and will not honor an IP address sent in across Connection 1, for example.
- 3) After establishing Connection 2, Domain A passes to Domain B two pieces of data: the token originally given to Domain A by Domain B, and a short-lived secret coined by Domain A.
- 4) When Domain B receives the secret over Connection 2 from Domain A, it sends the secret back to Domain A over Connection 1.
- 5) Upon validation of its secret, the Server at Domain A knows that the Connection 1 shares knowledge with Connection 2 and therefore can be assumed to exist between the same two points. Connection 1 and Connection 2, then can be trusted equally for the exchange of information.
- 6) If Connection 1 were trying to spoof Domain B, it would not receive Connection 2 from Domain A and would therefore not be able to obtain the necessary secret required to authenticate Connection 1.

Implementation Note: Domain B may want, at a minimum, to encode specific information about the requesting server, process, or thread as part of the token it passes to Domain A. In this way, Domain B can correlate Connection 1 with Connection 2.

6.3.2 Enhanced Security

The "Dial-Back" approach protects against all forms of attacks except man-in-the-middle attacks, packet snooping, and DNS hijack, and is therefore consistent with the objectives. Additionally, because many different End Users' Data will flow over a single connection, and a given End User's Data may be spread across multiple connections, the only way to actually spoof an End User is to effectively take over the entire domain. While that's certainly possible under this architecture, it's not very likely and would certainly be easily noticed.

If it became a requirement to protect against man-in-the-middle and packet snooping attacks, encrypted SSL connections can be established between the Servers without any change to the architecture.

If it became a requirement to protect against DNS hijacking, digital certificates with a known trust hierarchy could be used to authenticate Servers as part of the dial-back process.

Aoki & Wick

[Page 15]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

6.4 SPAM

When compared to Internet email, this proposal has the following benefits in fighting SPAM:

- 1) The source domain for an Instant Message is authenticated using the dial-back mechanism so it is not possible to easily spoof a domain. This doesn't mean it's not easy for a malicious entity to set up his own Instant Messaging Service and launch SPAM from it, but that domain will eventually get blocked as a SPAM domain.

It does mean, however, that one can't easily spoof messages from a high profile domain that one doesn't own.

- 2) There is no relaying in the protocol so there can be no accidental "open relay" servers for SPAMers to leverage.
- 3) Instant Messaging is generally not a store and forward system so a SPAMer would have to find the target online in order to send out a SPAM message. If a system supports offline Instant messages (which essentially creates a store and forward system), that system may need to implement some additional protections.
- 4) Early implementations should deploy rate limiting on all user accounts so any automated delivery mechanisms (e.g. agents that send out Instant Messages faster than a normal person would type) will get blocked. Exceptions may be made for approved notification services (e.g. an electronic trading company that lets End Users subscribe to stock price alerts), but the default will be that all accounts are rate limited.
- 5) Instant Messaging Clients may choose to automatically restrict or permit Instant Messages based on
 - Sender's domain
 - Sender's fully qualified address
 - Recipient's buddy or contact list
 - Other, vendor specific criteria

7. Authors' Addresses

Edwin Aoki
AOL
501 E. Middlefield Rd., MV-102
Mountain View, CA 94043
USA
aoki@aol.net

Aoki & Wick

[Page 16]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

Andy Wick
AOL
22000 AOL Way
Dulles, VA 20166
USA
wick@aol.net

8. Additional Documents

The following protocols and/or concepts described in this document will likely need additional explanation in additional documents:

- Detailed information about the IMX protocol
- The IMX DNS record
- The text/html-lite MIME type

9. Acknowledgements

Many thanks are due to the engineering and management team of America Online for their support in drafting this document. Particular credit is due to Eric Bosco and John Friend for their help in developing many of the ideas reflected in this paper.

10. References:

- [RFC 822] Crocker, D., "Standard for the Format of ARPA Internet Text Messages", STD 11, RFC 822, UDEL, August 1982.
- [RFC 2779] Day, M., Aggarwal, S., Mohr, G. and Vincent, J., "Instant Messaging / Presence Protocol Requirements", RFC 2779, February 2000.
- [RFC 2778] Day, M., Rosenberg, J. and H. Sagano, "A Model for Presence and Instant Messaging", RFC 2778, February 2000.
- [RFC 2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) - Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [RFC 1034] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, November 1987.

A. Appendix - Performance Against Objectives

In designing the specification for this architecture, the authors tried to strike a balance between practicality of implementation and the needs of End Users and Instant Messaging System operators. This section describes how, in the authors' view, the architecture described meets the requirements set out in Section 2.

Aoki & Wick

[Page 17]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

A.1 Scalability and Efficiency

It is obviously very important that servers implementing this architecture be able to scale to handle large numbers of users, messaging systems (domains) and requests across the protocol.

While avoiding the discussion of how to build a Server which can handle a large number of Client requests (i.e. users), the architecture itself should be scalable to large numbers of systems and efficient for heavy use, for the following reasons:

- 1) The authentication mechanism is lightweight, and authentication needs to occur only at the time a connection is established
- 2) Once authenticated, connections are basically stateless, so an implementation can easily use multiple hosts to split the load.

AOL 000361

- 3) There are a variety of ways to split the load for incoming connections across multiple hosts:
 - Multiple hosts can be specified in the IMX record
 - The load can be split using DNS round robin
 - The load can be split using a load balancing router
- 4) Because the architecture is Server to Server, Instant Messaging System vendors can optimize traffic within their own systems and use the interoperability architecture as needed for traffic across systems.
- 5) The protocol is very lightweight. The Gateway Servers are not being asked to do anything that's CPU intensive; mostly they are simply acting as a protocol translator between the public IMX protocol and the internal state of the Instant Messaging System.
- 6) The architecture works over semi-persistent connections that stay alive as long as there is traffic on them. This allows the traffic for many different users to flow over a single socket.
- 7) The protocol is envisioned to be very efficient and not "chatty". Each operation is self contained - fully encapsulating a given request in a single protocol request with a single protocol response.

[Page 18]

Aoki & Wick

INTERNET-DRAFT

IMX Architecture

June 15, 2000

A.2 Ease of Implementation

- 1) The architecture specifies a structure that allows vendors to quickly provide users with the benefit they most want - ability to message across systems, without precluding advanced features later.
- 2) There are three ways to participate in interoperability, allowing for all existing instant messaging products to choose the model which best suits its needs:
 - By having the system's Servers (the ones to which individual clients connect) be public and communicate with other systems,
 - By having each system's servers communicate through a Gateway which translates between a vendor's client-server protocol and the interoperability protocol, or
 - By having each system's Servers communicate via some private Server-to-Server protocol to an Aggregator, which translates between that protocol and the interoperability protocol.
- 3) It should be very straightforward to implement this protocol in a set of Gateway Servers that connect an existing system to the outside world, allowing vendors to leverage existing investments

AOL 000362

in technology and resources.

- 4) It is not anticipated that the Gateway Servers would need to hold any additional state beyond what a regular Instant Messaging System is already holding

A.2.3 Reliability

- 1) The stateless connections allow a given Server to go down, restart and resume processing of requests without a lengthy restart or connection building process.
- 2) The connect-on-demand capability makes for automatic recovery when a connection or Server goes down.

A.2.4 Security

- 1) Even without SSL, digital certificates, or end-to-end encryption the architecture provides for a system which is far more secure than e-mail.

Aoki & Wick

[Page 19]

INTERNET-DRAFT

IMX Architecture

June 15, 2000

- 2) Instant Messaging Service providers are responsible for the enforcement of security within their own domain, allowing them to use the most appropriate level of security for their user base.
- 3) Likely attacks on the system are difficult, noticeable, and can not easily target an individual account.
- 4) The architecture provides for vendors to implement more advanced security if appropriate.
- 5) The authenticated connection approach makes it more difficult for anonymous SPAM to enter the system.

Full Copyright Statement

"Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an

AOL 000363

"AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

Aoki & Wick

[Page 20]

AOL 000364

Open IM Architecture

